Supervenience and Infinitary Logic

Michael Glanzberg
Massachusetts Institute of Technology

The discussion of supervenience is replete with the use of infinitary logical operations. For instance, one may often find a supervenient property that corresponds to an infinite collection of supervenience-base properties, and then ask about the infinite disjunction of all those base properties. This is crucial to a well-known argument of Kim (1984) that supervenience comes nearer to reduction than many non-reductive physicalists suppose. It also appears in recent discussions such as Jackson (1998).

Some philosophers have been troubled simply by the infinity of such a disjunction. Logicians tend to react somewhat differently. Infinitary logical operations have been studies in depth by the highly developed field of infinitary logic, and many of their properties are well-understood. On the other hand, as anyone who has ever worked with infinitary logic knows, it has proved difficult, messy, and filled with surprising pitfalls. Moreover, there are lots of different infinitary logics, displaying different characteristics, some better, some worse. Logicians are not likely to object to infinitary operations per se, but they are likely to ask whether their application to a metaphysical issue like supervenience work as smoothly as it is sometimes assumed.

In this paper, I shall investigate the interaction between supervenience and infinitary logic. Supervenience has long been a point of contact between logic and metaphysics. In examining it, some philosophers have already questioned the metaphysical status of certain finitary logical operations. I shall show here that the step to infinitary logical operations raises significant metaphysical issues of its own. This step, I shall argue, is not an all-or-nothing deal. If we accept the use of infinitary logical operations, we still face hard choices about the strength of such operations to allow. This in turn forces us to confront difficult metaphysical questions about what to count as a property of a given class. In this respect, I shall argue, the step to infinitary operations leads to more complex and far-reaching versions of problems we encountered with finitary operations. The step to the infinitary is by no means banal, even for those who take the finitary logical operations for granted. I shall go on to argue that when we look at the underlying source of these problems, we can in fact see that the step to the infinitary raises some special metaphysical problems of its own, different from those that appear with its finitary counterpart. In particular, it raises a fundamental question about what counts as the complete physical state of the universe.

The arguments of this paper are divided into five sections. Section (1) sets up the framework for the application of infinitary logic to issues of supervenience. Section (2) provides an number of examples of surprising and unwelcome consequences of the unrestricted use of infinitary logic. These provide good reason to look for ways to restrict it somehow. In Section (3), a diagnosis is offered of the source of the problems raised in Section (2). This leads to a discussion of how infinitary logic may be appropriately restricted in Section (4). Once we see how to restrict infinitary logic, we will see that choosing among restrictions leads to serious metaphysical questions. The ones that appear in Section (4) are refined and extended versions of problems that have appeared with respect to finitary logic. In section (5), I show that considering infinitary logic leads to a puzzle about the complete physical state of the universe. This provides a metaphysical issue raised by infinitary logic that does not have an analog in the finitary.

1. Preliminaries: Supervenience, Closure, and Infinitary Logic

Supervenience relations are relations between classes of properties. A supervenience relation holds between a class of properties $\mathcal{A}$, the supervening class, and a class of properties $\mathcal{B}$, the base class, if whenever there is sameness of $\mathcal{B}$-properties, there must be sameness of $\mathcal{A}$-properties.[1]

This is not a definition, but a general scheme into which many different definitions may fit. Many varieties of supervenience have been identified: weak, strong, global, regional, and others. They differ along two axes. First, how to spell out the idea of sameness or difference in a class of properties. Second, what modal force to provide. However, for purposes of this paper, these issues will be secondary. It will hence be useful to fix the core idea of which these are elaborations. Write 'x has the same C-properties as y' as '$x \sim_C y$'. The the core notion is simply covariation of $\mathcal{A}$ with $\mathcal{B}$, which we may express as:

$$\forall x \forall y (x \sim_{\mathcal{B}} y \rightarrow x \sim_{\mathcal{A}} y).$$

To get supervenience from this, one needs to fix what x and y are. Options include individual objects, n-tuples of objects, regions or parts of possible worlds, or entire possible worlds. One also needs to appropriately modalize covariation.[2]

Once we focus on covariation, it is easy to see the importance of what closure assumptions we may make about a given class of properties. The stronger the closure assumptions we make about the base class $\mathcal{B}$, the harder it is to have sameness of $\mathcal{B}$, and hence the easier it is to have covariation. In an extreme case, closure assumptions can trivialize covariation, and hence supervenience, entirely. If the closure assumptions on $\mathcal{B}$ imply that there are enough $\mathcal{B}$-properties that we never have sameness of $\mathcal{B}$, for instance, then covariation of $\mathcal{A}$ with $\mathcal{B}$ is clearly trivial. These same holds if the closure conditions imply that $\mathcal{A} \subseteq \mathcal{B}$.

The closure conditions that have been discussed in the literature are usually not so absurd. Most attention has been paid to logical closure conditions, especially closure under Boolean operations. Closure under Boolean operations is at the very least useful. For instance, it allows some common variants of familiar supervenience relations to be shown equivalent.[3] It is also controversial. Armstrong (1978), for instances, questions whether negation and disjunction are always property-forming operations at all. With an eye more towards the properties of a physical supervenience base, closure under disjunction has been questioned by Teller (1983a). Post (1983) and Van Cleve (1990) have challenged closure under negation. On the other hand, Van Cleve defends closure under disjunction. Kim mounts a spirited defense of Boolean closure, saying its rejection "would work havoc on with free and creative scientific theorizing." (Kim 1990, p. 153.)

I am somewhat sympathetic to Boolean closure, but it is not my purpose to argue the point here. Rather, my concern is with closure under infinitary Boolean operations. In order to make some more refined distinctions about infinitary closure conditions, it will be helpful to reformulate them in the terms of formal logic.

For a class C of properties, start with a collection of predicate and relation symbols corresponding to each (primitive) element of C (for most applications, the supervenience base class $\mathcal{B}$). Let $\mathcal{L}$ be the usual first order language whose sentences are built out of these predicate and relation symbols using the logical operators $\wedge$, $\vee$, $\neg$, $\rightarrow$, $\forall$, $\exists$, $=$ (or anything else your favorite syntax might use). Sometimes, I shall assume that $\mathcal{L}$ contains constants to name any object. However, the interaction of names and identify raises some difficulties that we will have to keep track of below, so unless otherwise specified, assume $\mathcal{L}$ does not contain constants. Note that $\mathcal{L}$ is an interpreted language. In fact its interpretation is by intensional elements, like properties, but for most of what follows we may think of it as an ordinary interpreted extensional first order language.

Now, in place of Boolean closure of the class C, we may talk of $\mathcal{L}$-closure: the assumption that any formula of $\mathcal{L}$ corresponds to an element of C. We only really care about those formulas with at least one free variable, but throwing elements corresponding to sentences—something more like facts than properties—changes nothing important here. If the class C has no relations, the step from Boolean to $\mathcal{L}$-closure is vacuous. If there are relations, then it is no more than a technical convenience. If we have a relation $R(x, y)$ in C, we assume that we also have in C such properties as $\exists y R(x, y)$. If we are to work with relations, we need to have these available. Note that if we have a name 'a', then we also have $R(x, a)$.[4]

To describe infinitary closure, we simply modify our syntax to allow disjunctions and conjunctions of infinite sets of formulas. For any set of formulas $\Phi = \{\phi_1, \phi_2, \ldots\}$, there is a formula $\bigvee \Phi$ which represents the infinite $\phi_1 \vee \phi_2 \ldots$ and $\bigwedge \Phi$ which represents $\phi_1 \wedge \phi_2 \ldots$. ($\Phi$ does not need to be an ordered set. I am just pretending it is for illustrative purposes.) To simplify notation, I shall sometimes write $\bigvee\{\phi_i \mid i \in I\}$ as $\bigvee_{i \in I} \phi_i$. If we allow only finite quantifier prefixes, we have the language logicians call $\mathcal{L}_{\infty\omega}$. For formal purposes, in place of infinitary Boolean closure, we may talk of $\mathcal{L}_{\infty\omega}$-closure.[5]

Infinitary closure has certainly raised some eyebrows,[6] but for the most part, those who accept Boolean closure at all have taken infinitary closure for granted. Each of the authors cited in Note (3) in fact supposes infinitary closure. The general assumption seems to be that infinitary closure stands or falls with finitary Boolean closure. It seems to be assumed that this is an unremarkable technical point, or perhaps a matter of logic. Something of this is echoed in Kim's defense of infinitary closure:

> I don't see any special problem with an infinite procedure here, any more than in the case of forming infinite unions of sets or the addition of infinite series of numbers. (Kim 1990, p. 152.)

As a technical point, the general assumption is a mistake. $\mathcal{L}_{\infty\omega}$ behaves very differently from $\mathcal{L}$, and one could find all kinds of reasons for rejecting full $\mathcal{L}_{\infty\omega}$-closure that have nothing to do with any worries about the coherence of infinitary procedures. I shall provide some examples related to supervenience in Section (2). More importantly, those inclined to accept some infinitary closure are not thereby committed to accepting all of $\mathcal{L}_{\infty\omega}$-closure. Between $\mathcal{L}$ and $\mathcal{L}_{\infty\omega}$ is a rich range of logics, with substantially different properties, as I shall discuss in Section (4).

My primary aim here is not to challenge or defend Boolean closure. Rather, I shall argue that those inclined to accept some infinitary Boolean closure face some difficult choices in deciding how much infinitary closure to accept. Moreover, these choices are not mere technicalities; they amount to some hard metaphysical issues. Some of these, we shall see, are distinct from those raised by the question of finitary closure.

2. The Perils of Infinitary Closure

In this section, I shall compile a list of troubling consequences of $\mathcal{L}_{\infty\omega}$-closure. Different views will no doubt find some of them more troubling than others, but taken together, they provide a preponderance of evidence full $\mathcal{L}_{\infty\omega}$-closure is unacceptable. Thus, especially if we are inclined to accept some infinitary closure, we have good reason to look to infinitary languages which substantially restrict $\mathcal{L}_{\infty\omega}$. It is in deciding how to restrict infinitary closure that we encounter metaphysical problems, as I shall argue in sections to follow.

A. The Reduction Argument. Kim (1984) argues that strong supervenience implies reduction. Strong supervenience is based on individualistic covariation. In the covariation schema above, x and y are taken to be individual objects. As has been much discussed, strong supervenience modalizes covariation to give it the force of inter-world comparison. The use of infinitary closure, however, really only interacts with covariation. To make this clear, we may divide the argument into three steps.

Step 1: disjunctive $\mathcal{B}$-surrogates. The idea here, roughly, is to find a way to replace $\mathcal{A}$-talk with $\mathcal{B}$-talk, by finding a surrogate for each $\mathcal{A}$-property within $\mathcal{B}$.

Pick any $A \in \mathcal{A}$. To build a surrogate for A, we start by taking any x such that $A(x)$. We then look at the collection $\{B \in \mathcal{B} \mid B(x)\}$ of all the properties in $\mathcal{B}$ which are had by x. Let $\beta_x = \bigwedge\{B \in \mathcal{B} \mid B(x)\}$. To build the $\mathcal{B}$-surrogate for A, we simply form the disjunction $\bigvee\{\beta_x \mid A(x)\}$. Call this '$\alpha$'. (For uninstantiated A, pick or build an uninstantiated member of $\mathcal{B}$ any way you like.)

It is easy to see that covariation implies $\forall x(A(x) \leftrightarrow \alpha(x))$. Furthermore, $\mathcal{L}_{\infty\omega}$-closure ensures that $\alpha \in \mathcal{B}$. Hence, for each A in $\mathcal{A}$, we have a surrogate α in $\mathcal{B}$.

Step 2: necessary coextensiveness. Once we have identified α as a surrogate for A, it can be shown that strong supervenience implies that α and A are not merely coextensive, but necessarily coextensive.[7]

Step 3: Reduction (?). The final step is to argue that the availability for each $\mathcal{A}$-property of a necessarily coextensive surrogate $\mathcal{B}$-property amounts to the reduction of $\mathcal{A}$ to $\mathcal{B}$.

Some, such as Jackson (1998), take this conclusion to be evident. Others, such as Kim, are more cautious. It is striking, as Kim notes, that the necessary equivalence established in step (2) provides a 'bridge-law', which implies that by the standards of the deductive-nomological model of reduction, $\mathcal{A}$-properties, or theories of them, can be reduced to $\mathcal{B}$-properties or theories. On the other hand, the deductive-nomological model is by no means universally accepted. Whether or not by some other more nuanced standard the construction of steps (1) and (2) is sufficient for reduction is, of course, a subtle matter. Kim himself is willing to conclude that at least in a "somewhat attenuated sense" (1990, p. 154) it is.

Whether or not it is genuinely reduction or something less, the presence of a necessarily equivalent surrogate in $\mathcal{B}$ for every property in $\mathcal{A}$ is perhaps the most well-known consequence of $\mathcal{L}_{\infty\omega}$-closure. Though well-known, it is not the most threatening of its consequences, as we will now see.

B. Fully Describing a Possible World. The next consequence of $\mathcal{L}_{\infty\omega}$-closure is that $\mathcal{B}$ contains enough information to fully describe possible worlds. This is surprising enough on its own, and it will be useful for generating other results as well.

Let us think of each possible world as given by a structure. (Specifically, an $\mathcal{L}$-structure in the usual logician's sense of a domain for quantifiers and extensions for all the relations of $\mathcal{L}$.) This is the usual assumption of modal logic, and as a matter of technical convenience, it is unremarkable.

So, fix some structure $\mathfrak{M}$ with domain M. If we add constants $\dot{m}$ for each element m of the domain M, it is entirely easy to to use $\mathcal{L}_{\infty\omega}$ to describe $\mathfrak{M}$ completely, i.e. up to isomorphism. Let $Diag(\mathfrak{M})$ be the diagram of $\mathfrak{M}$: the collection of atomic and negated atomic sentences true in $\mathfrak{M}$. Then $\mathfrak{M}$ is characterized up to isomorphism by:

$$\bigwedge Diag(\mathfrak{M}) \wedge \forall x(\bigvee_{m\in M} x = \dot{m}).$$

There may well be ground for worrying about the use of the predicates $x = \dot{m}$. However, it is one of the really deep results about $\mathcal{L}_{\infty\omega}$ that they are not needed. Even if $\mathcal{L}$ contains no names, for any $\mathcal{L}$-structure $\mathfrak{M}$ there is an $\mathcal{L}_{\infty\omega}$-sentence $\sigma_{\mathfrak{M}}$—called the Scott sentence of $\mathfrak{M}$—which fully describes $\mathfrak{M}$.[8]

C. Resplicing. One of the more controversial closure conditions to be discussed in the literature is resplicing: for any property P such that its extension P_w in world w is the extension of some property in $\mathcal{B}$, $P \in \mathcal{B}$. This makes $\mathcal{B}$ closed under arbitrary rearrangements of extensions of its properties. This condition has the surprising effect of making strong and weak supervenience equivalent. Many have found this untenable.[9]

One of the most striking consequences of $\mathcal{L}_{\infty\omega}$-closure is that it implies closure under resplicing. This can be shown by using our technique for fully describing a possible world. Let W be the set of worlds. From our construction above, we have for each $w \in W$ a sentence σ_w which describes it. Now, pick some set of extensions of $\mathcal{B}$-properties $\{E_w \mid w \in W\}$ to be respliced. Each E_w is the extension of some $\mathcal{L}_{\infty\omega}$-formula ϕ_w in w. Then, the resplicing of this set is given by:

$$\bigwedge_{w \in W} (\sigma_w \to \phi_w).$$

Under $\mathcal{L}_{\infty\omega}$-closure, this is a property in $\mathcal{B}$, so $\mathcal{B}$ is closed under resplicing.

D. Adding an Arbitrary Set. Resplicing has been challenged for its consequences for strong and weak supervenience, and simply for the implausibility of the properties it builds. In many cases, $\mathcal{L}_{\infty\omega}$-closure can in a much more direct way add all sorts of unlikely properties to $\mathcal{B}$.

We begin by looking at sets. $\mathcal{L}_{\infty\omega}$ can in some cases define an arbitrary set. I shall provide two examples, which have different metaphysical costs.

Pick some set of objects, say P. P may be as arbitrary, gerrymandered, or wild as you like. Suppose, first, that we have either names $\dot{p}$ for every element of P, or that we can add them. Then, using these names and identity, we can describe P in $\mathcal{L}_{\infty\omega}$ simply as:

$$\bigvee_{p \in P} x = \dot{p}.$$

This construction is entirely independent of the class $\mathcal{B}$. Hence, it shows that $\mathcal{L}_{\infty\omega}$-closure of any class implies that it contains any set of objects that might be named.

It is crucial to this construction that we use the predicates $x = \dot{p}$, and this is no doubt controversial. For one thing, if we assume we already have names for every object, and allow the properties $x = \dot{p}$, we trivialize supervenience. However, we do not need to assume this, but only that we can introduce names for some collection of objects. I do believe there is a variety of physicalism which should accept this for a physical supervenience base. After all, naming objects is certainly something which goes on in the physics lab all the time. Even so, many current physicalists will not accept it. For instance, Lewis (1986) explicitly require the supervenience base to be purely qualitative, which surely disallows anything like $x = \dot{p}$. Others will object that these amount to adding essences, or haecceities, which already go beyond the physical.[10]

There are, however, ways to do the same thing without name and identity, in some cases. All we need is a world where enough combinations of properties allow us to narrow down to a single object. To take a simple example, suppose our world has objects arranged like the rational numbers, and properties corresponding to open intervals. Now, each such interval is filled with lots of rational numbers. We are to start with nowhere close to uniquely identifying properties. We have nothing that appears to go beyond the qualitative. But, we can narrow down to single rational numbers by taking intersections of intervals:

$$\frac{n}{m} = \bigcap\{(a, b) \mid \frac{n}{m} \in (a, b)\}.$$

Now, I used the name $\frac{n}{m}$ to indicate this, but it is just the intersection of some collection of intervals, say $\bigcap \Phi$.

Each interval corresponds to a predicate of $\mathcal{L}$. Hence, this intersection is in fact an $\mathcal{L}_{\infty\omega}$-construction. We have:

$$\Xi(x) \leftrightarrow \bigwedge \Phi(x),$$

where each member of Φ is a predicate for an interval. This relies on no names and no identity—nothing that appears to be non-qualitative. Yet Ξ holds of exactly one object: $\frac{n}{m}$. We can form a corresponding predicate Ξ_p for any object (rational number) p. We can thus construct any set of objects P as:

$$\bigvee_{p \in P} \Xi_p,$$

more or less as we did above. Indeed, we may have more to worry about, as with the family of Ξ_p we have uniquely individuating properties for all objects, which is a disaster for supervenience.

Now, exactly how we build the Ξs will depend on just how objects and properties are distributed in a world, and there may well be some worlds where we cannot do it. However, closure assumptions are not the sort of thing we choose to invoke in a given world, and then suspend in others. So the fact that in some worlds we can construct the $\Xi\underline{s}$ using $\mathcal{L}_{\infty\omega}$ is yet more grounds for worry about $\mathcal{L}_{\infty\omega}$-closure.

E. Adding an Arbitrary Property. We now have $\mathcal{L}_{\infty\omega}$-constructions which describe arbitrary sets and arbitrary worlds. Together, these allow us to add more or less arbitrary properties to any class of properties. Pick some set P_w for each world w. Once we have $\mathcal{L}_{\infty\omega}$-formulas ϕ_w for each P_w, and σ_w for each world w, we can give the property with extension P_w in world w by:

$$\bigwedge_{w}(\sigma_w \to \phi_w).$$

In a modal context, we can think of sets as rigid properties: those with constant extension. $\mathcal{L}_{\infty\omega}$-closure can have the result that any such property is in $\mathcal{B}$, as well as many arbitrary non-rigid properties as well. Just looking at the rigid properties provides a startling result. A physicalist, for instance, would have to grant that the base class $\mathcal{B}$ contains rigid properties for any arbitrary set of objects, including objects whose only feature in common is some moral property, or aesthetic property, or nothing more than appearing on some random list. The substantiality of supervenience is clearly threatened.[11]

I conclude that the assumption of full $\mathcal{L}_{\infty\omega}$-closure has too many unacceptable consequences, and must be restricted in some way. As I mentioned, I do not insist that everyone find all the consequences odious. Opinions differ as to which are problems, and how bad. However, the preponderance of evidence is that $\mathcal{L}_{\infty\omega}$ is too much.

3. The Source of the Problems

Before looking at the technical resources for restricting $\mathcal{L}_{\infty\omega}$, we should ask what really lies behind the examples of the last section. The common symptom of all of them is unexpected growth of the base class. It grows to contain disjunctive surrogates in the reduction argument, or resplices, or even arbitrary sets and properties.

The example of adding an arbitrary set, I believe, shows must clearly what the cause of this symptom is. $\mathcal{L}_{\infty\omega}$-closure enables us to carry out essentially set-theoretic constructions within a class of properties. In the example of adding a set, we do set theory by big disjunctions and conjunctions of claims about the members of sets, but we do some set theory nonetheless. Saying $(\bigvee_{p \in P} x = \dot{p})(a)$, for instance, is just to say $a \in P$. Though it is less blatant, set theory is at work in the other examples as well. We do some more subtle set theory to provide reconstructions of possible worlds, and then use them to combine other sets into resplicings. The crucial step (1) of the reduction argument can be thought of in terms of sets of base properties as well as in terms of disjunctions. Assuming full $\mathcal{L}_{\infty\omega}$-closure amounts to assuming our class $\mathcal{B}$ is closed under a broad range of set-theoretic constructions.[12]

We should well expect a closure assumption that builds in this much set theory to have all sorts of untoward consequences. Set theory can hardly be expected to respect physicalist scruples, or any other requirements we might place on the base class $\mathcal{B}$. Quite the reverse, set theory has a way of smuggling in all kinds of things into $\mathcal{B}$, whether we think they belong there or not. Sometimes, as in the addition of arbitrary properties case, set theory simply puts things in $\mathcal{B}$ that have nothing to do with what started out in it. It just throws them in. Other times, as in the reduction argument case, set theory enables us to build new properties out of what we already had in $\mathcal{B}$. When these combine, the result is that the class $\mathcal{B}$ is expanded well beyond what we can tolerate.

It will come as little surprise that with enough set theory, one can build at least surrogates for all sorts of properties, including the wildly arbitrary ones that we considered above. If the reduction argument had claimed only that in set theory, surrogates for supervening properties can be constructed, I doubt it would have raised many eyebrows, or been of much concern to the non-reductive physicalists at which it takes aim. Set theory is, of course, a theory—one very well-suited to this kind of surrogate construction. The surprise of the construction came because it looked like it did not use anything like set theory. It seemed to use only some innocuous forms of reasoning—only logic. Infinitary logic does not look like a theory; it looks like logic.

The moral of the discussion above is not that infinitary logic is bad, but rather that one must not let looks deceive. Assuming infinitary closure is to assume some degree of closure under set-theoretic constructions. One must decide if and how much set-theory-like construction one is going to include in the closure conditions on the base class. Thus, in spite of the way things look, one must decide how much infinitary logic to allow.

We will see in the next section that there are reasonable ways to do this. They will require us to make some very fine-grained choices among sets. As these choices amount to deciding what is in a metaphysically significant supervenience base class, we can already see that they will force us to confront some metaphysical issues.

4. What to Do

Unrestricted $\mathcal{L}_{\infty\omega}$-closure amounts to closing under a significant amount of set theory. The examples of Section (2) show that this can have some unwelcome consequences. Of course, those who reject Boolean closure in any form will find comfort in our dismay. Notwithstanding, the question I shall pursue is whether there is some way to keep Boolean closure, grant to Kim and others the point that once we accept Boolean closure, simply refusing to consider any infinitary Boolean operations appears poorly motivated and difficult to defend, and yet rein in the amount of set theory we build in along the way.

There is a way. The task is to find some suitably fine-grained way to restrict $\mathcal{L}_{\infty\omega}$ without cutting all the way back to the finitary $\mathcal{L}$. It turns out there is a very elegant way to do so. The basic construction that makes infinitary logic infinitary is $\bigvee\Phi$. $\mathcal{L}_{\infty\omega}$ places no limits on which sets Φ can be used. What is needed is some principled way to place limits on which sets Φ are allowed. Experience has shown that the most effective way to do so is based in their <u>complexity</u>.[13]

The idea is to allow an infinite disjunction $\bigvee\Phi$ only if Φ is a reasonably simple, reasonably 'nice', set. The way this is done is to look for a mini-universe of 'nice' sets, and to allow an infinitary conjunction or disjunction only if the set of formulas conjoined or disjoined can be thought of as coded up as a set in the mini-universe of nice sets. What counts as nice? Really nice is finite, where everything works out simply. So nice is having enough features of the finite case, taken loosely enough to allow a rich classification of infinite sets by how nice they are.

It turns out there is a very natural definition of 'nice' in this sense. A mini-universe of nice sets is best identified with what is technically known as an admissible set. I shall not formally define this notion, as it is an entirely technical matter. I shall mention one important feature of admissible sets, which might help to give a sense of why they are useful. Given an admissible set A, there is a next biggest admissible set A^+. Given a certain degree of niceness A, we can identify the next-less-nice collection A^+. The progression $A, A^+, A^{++}, \ldots$ produces a kind of ranking of how complex the elements of these collections are. It also turns out that admissibles have important connections to definability theory. The question of where in the progression a set falls comes down to how complex its definition needs to be.[14]

One of the truly remarkable ideas in the study of infinitary logic, due to Jon Barwise, is to apply the notion of admissible (nice) set to it. The basic idea, as I said, is to insist that when we disjoin or conjoin a set of formulas Φ, the set be nice. Actually, it turns out the right idea is to insist that all formulas be nice. We must think of each formula of infinitary logic as coded by a set. $\bigvee \Phi$ is the set $\langle \vee, \Phi \rangle$. Admissible sets are closed under pairing, so to require that Φ be nice is just to require that $\langle \vee, \Phi \rangle$ be nice.

The right way to restrict infinitary logic, it turns out, is to require that all formulas be nice: we find some admissible set, and restrict infinitary logic to formulas that are in it. Technically, this gives what is called an admissible fragment of infinitary logic.[15] The fragment produced by an admissible set A is usually written $\mathcal{L}_A$. Countable admissible fragments turn out to be technically very nice indeed. They behave very much like first-order logic, especially in yielding strong completeness and compactness theorems.

Restricting infinitary logic this way provides a technical basis for responding to the kinds of problems we saw in Sections (2) and (3). Whether or not an infinitary operation like $\bigvee \Phi$ is allowed depends on the set Φ. For a given admissible fragment $\mathcal{L}_A$, it is allowed only if the set Φ itself is in A—is niece enough. As a result, the set-theoretic constructions that can be carried out in $\mathcal{L}_A$ are restricted to those that fall in A. For example, the formula $\bigvee_{p \in P} x = \dot{p}$ is in a given $\mathcal{L}_A$ just in case the set P is in A. Hence, the question of whether or not we allow this infinitary formula comes down to the question of whether or not we include the set it defines. We can thus use the restriction to a suitable $\mathcal{L}_A$ to block the construction for adding an arbitrary set of Section (2). The same goes for the construction for adding an arbitrary property. Likewise, in the reduction argument, the formula $\bigvee\{\beta_x \mid A(x)\}$ is only allowed if $\{\beta_x \mid A(x)\}$ is in A. Each problem of Section (2) can be avoided by restricting to a sufficiently narrow $\mathcal{L}_A$, and thereby limiting the set-theoretic constructions we allow for the class $\mathcal{B}$.[16]

If we want to have infinitary closure, we have to decide how much. We may decide by choosing an admissible fragment $\mathcal{L}_A$ to describe closure. This will involve deciding whether certain formulas are acceptable or not. Many of these formulas describe sets or properties—they embody set-theoretic operations. In choosing how much closure to have, one will have to decide which sets and properties to allow.

The choices must be very detailed. There are lots of admissible sets, and so lots of admissible fragments of infinitary logic. In fact, for any infinite set, there are admissibles that it contain it, and admissibles that do not.[17] The result is that in choosing how much infinitary closure to allow, we will have to make specific choices about whether or not to accept each of a huge range of sets and properties.

We can now see clearly that this is not simply a matter of whether infinitary closure in general is coherent or acceptable. Even if we grant that it may be, we have to decide how much infinitary logic—which fragments—to accept. There are enough admissible fragments available to block any of the consequences of Section (2). Indeed, there are so many that one can pick and choose among these consequences pretty much at will. But the cost is that one also has to pick and choose equally among sets and properties.

Of course, we had to be prepared to make decisions of this kind. Anyone interested in a class of properties for some philosophical purpose must be prepared to say what should be in it. The question of finitary closure already raised the issue of what ways elements of the class can be combined. At the very least, I have shown that these questions appear in the extension from finitary to infinitary closure as well. Those who think that infinitary closure comes for free are mistaken. Indeed, I believe we have shown somewhat more. The range of choices that must be made in the infinitary case is more extensive. For any given set, property, or collection of properties, there will be some admissibles—some infinitary closure conditions—that include it, and some that do not. For each infinite set or set-theoretically describable way of building properties, we will face a distinct choice. We will have to make far more discriminating choices about closure than just, say, whether we like negation or disjunction. The technically available range of options in the infinitary case is so rich that it forces an equally rich range of metaphysical decisions on us.

The sort of choice we face in the infinitary case is well-illustrated by the reduction argument. We are offered a property $\bigvee\{\beta_x \mid A(x)\}$. We can rule this in or out, depending on whether we decide that the collection $\{\beta_x \mid A(x)\}$ is a suitable base for building a disjunctive property. Technically, we can find admissibles that contain this set, and ones that do not (assuming it is infinite), so either option is open. One might wish to argue that it should be ruled in. For instance, one might argue the $\beta_x\underline{s}$ are not just some arbitrary collection, but rather they are already brought together by $A(x)$, and so are a set we should rule in. Others will reject this, on grounds that $A(x)$ is merely supervenient. Either way, my point is, the argument (for those not already determined to reject all infinitary disjunctions) comes down to some such specific question. We will face an equally specific question for any of the rich range of set-theoretic constructions that can be carried out in $\mathcal{L}_{\infty\omega}$. Taken together, these will require us to make a very wide range of highly detailed metaphysical choices: wider and more detailed than we faced in the finitary case. Far from being a mere technicality, the technical situation makes the step from finitary to infinitary closure one that confronts us with much wider and much more subtle range of metaphysical issues than we had before.

5. The State of Things

So far, I have argued that infinitary closure raises a full range of subtle metaphysical problems. I suggested in the last section that we will face more, and more specific, choices than we did in the case of finitary closure. Hence, I suggested, the step to infinitary closure makes for a more complex metaphysical situation. Yet so far, the most I have shown is that the step to the infinitary raises additional complex questions of a kind we were already prepared to answer. We were prepared to decide what to count as a supervenience base property, and which combinations of base properties to allow. I have argued the step to the infinitary raises a surprising array of additional complexities, but they remain issues of this sort.

This is not the end of the story. I shall now argue that the underlying source of the complexities we have seen in fact raises a metaphysical puzzle all its own.

To see this, we must return to the basic idea behind supervenience. As has been much remarked, supervenience is supposed to capture a sort of determination. A physicalist about the mental, for instance, might say that the physical properties of a person (her brain? her environment?) determines her mental properties. Many want to capture this situation by an appropriate supervenience relation between mental and physical properties. This is sometimes explained by way of the metaphor of the Laplacian demon (as in Horgan 1983). The demon fixes the base facts, by distributing base properties among appropriate objects. The idea of determination is that the demon has no further task to do to fix the supervening facts.

The special problem infinitary logic raises is one of what counts as the demon's tasks. The demon is to fix all the base facts—fix the base state of the universe. An initial assessment of these tasks seems to be that they are described by the true finitary sentences that can be expressed in the appropriate language—all the true $\mathcal{L}$-sentences. We will see, however, that if we expand our horizons to the infinitary, we can find additional facts corresponding to infinitary sentences, that are not on our list of the demon's tasks. Yet these are expressed in the language of the base, so it appears we have found additional base facts. We thought we had a complete list of the demon's tasks—the complete base state of the universe—but now we appear to have found something left out. As a result, we shall see, infinitary logic makes trouble for the idea that there is a complete state of the universe, relative to some class of properties. We face a problem of how to make sense of the complete physical state of the universe, for example.

Let me explain a little more fully. We know that in many cases, there is no single finitary sentence expressing the state of the universe. If there are infinitely many things to say—infinitely many combinations of properties and objects bearing them—there is none. We might have thought that this is not so for infinitary logic. After all, in infinitary logic we can conjoin all the true finitary sentences. Is this not the statement of the complete state? If we conjoin all the true finitary sentences of the language suitable for expressing physical properties, for instance, have we not expressed the complete physical state of the universe?

In general, no. Once we grant that we can express things using an infinitary language, we no longer have any guarantee it is. In many cases, we can show that there is something true which is expressible only in the infinitary language, and so left out of the conjunction of true finitary sentences. Insofar as the infinitary language is still the language of, say, physical properties, we may well have to conclude there is some aspect of the physical state of the universe that we failed to capture in our attempt to express its complete physical state.

This is not an entirely trivial result. We might have thought that though infinitary logic gives us more sentences, there is nothing really new to say. We might have expected infinitary sentences just to repeat at length what we said with combinations of finitary sentences before. But the behavior we saw above should make us doubt this, and indeed it is not so. It is a fact that many of the nicest infinitary languages cannot express the complete state of things relative to that language (the conjunction of all true sentences of that language). The proof of the general fact involves a couple of technical difficulties, and so is consigned to an appendix.

Without going into the details now, what we have already seen should give some sense of why this is so. The problems raised in Section (2) and the diagnosis of Section (3) all point to what we might call the expansionist tendency of infinitary logic. Given some resources, like some class of properties, infinitary logic tends to allow one to exceed that class. We have seen how it can do so by providing set-theoretic constructions which can be used to build new elements not in the class with which we started. If we start with a class of true sentences, infinitary logic allows us to find a new true sentence not in the class. We can find one that is genuinely new, in that it expresses something different from anything in the initial class. Hence, if we start with some list of sentence which we take to provide the complete state of the universe, infinitary logic provides us with a new sentence not on the list. It shows there had to be something we left out.

This creates a puzzle for what the demon's tasks are. Initially, it appeared they were more than adequately described by all the true finitary sentences. But once we allow infinitary logic, we can find new facts, not on the demon's list. These are expressed in the infinitary $\mathcal{L}_{\infty\omega}$; yet insofar as they are expressed in the vocabulary of the base class, it appears they should have been on the demon's list after all. They appear to be perfectly good base facts. The situation now worsens, if we have all of $\mathcal{L}_{\infty\omega}$ available. It is no help simply to grant we may have missed something, and throw in the new fact. The result is yet again a set of sentences of $\mathcal{L}_{\infty\omega}$, and so again we will be able to find a further new fact, expressed by a sentence not on the list. If we start this process, and have all of $\mathcal{L}_{\infty\omega}$, we can never finish. We can never come to the complete state of the universe.

This puzzle has two aspects. First, it is a reflection of the issues we have been discussing all along. The additional facts the process generates have just the features we expect of supervenient facts. They are not fixed directly by the demon, in that they are not the (initial!) list of its tasks to perform. Yet they are determined by what the demon does. The truth of the infinitary sentences is determined by the truth of the finitary ones.[18] It is striking, though, how close they are to the base facts the demon does fix. They are not the sorts of facts, like mental ones, which seem to be expressible only in some other vocabulary. They are expressed in the same vocabulary, with only some extra logical resources. For this reason, we might have been inclined to count them as indeed part of the base. The assumption of $\mathcal{L}_{\infty\omega}$-closure does so.

The temptation here is to throw lots of infinitely complex properties into the base. The discussion of the preceding sections shows that this is a temptation to go down a dangerous path. The reduction argument shows that if we throw enough complex properties into the base, we wind up with surrogates for all the supervening properties, whether they looked like they were just infinitely complex base properties or not. As we have seen, there are much more dire consequences in store. We have seen it is easy to throw in so many complex properties as to weaken the supervenience base to the point of trivialization.

In Section (4), we saw how to avoid disaster by restricting closure to some suitable $\mathcal{L}_A$. As we reflected on the constructions of Section (2), this appeared to be the correct approach, though it raised a range of complex metaphysical problems. I believe it is the correct approach, but now, when we apply it to the puzzle of the demon's tasks, we encounter a more perplexing problem. This is the second aspect of the puzzle.

We may indeed apply the solution of Section (4) to the demon's tasks. We may say that all the base facts there are are those that correspond to true sentences of some suitably chosen $\mathcal{L}_A$, and thereby fix the list of the demon's tasks—thereby fix what we will count as the complete base state of the universe. We can indeed construct a new true sentence of $\mathcal{L}_{\infty\omega}$ not on the list, but it will not normally be a sentence of $\mathcal{L}_A$, so we may simply insist that it does not give us a new base fact we left out. We might, for instance, say that $\mathcal{L}_A$ gives us all the physical facts, and while the new sentence gives us a fact, it is not a physical fact.[19]

The problem is that it is very hard to find any reason to stop counting the new sentences as giving physical facts. In the cases of Section (2), we might well have found good metaphysical reason to rule out some property as too gerrymandered to be physical, for instance. But nothing like this is going on in the case of the demon's list. All we need to do to find a new true sentence which goes beyond the list is to collect together the sentences on the list and say 'all of those'. This will hardly be gerrymandered, or otherwise metaphysically frivolous. Indeed, even outright resistance to any infinitary closure of any kind does not seem to provide a satisfactory answer. When talking about the complete state of the universe, we are already talking about something infinitary. The demon's list of tasks is infinite. To fix the complete state of the universe, the demon must already perform an infinite supertask. If we then find that in the realm of such infinite supertasks there is something we left off the list, one more infinitely complex arrangement of physical properties, what grounds have we to insist it is not a further physical property?

To get a further sense of why this problem is hard, it is useful to shift our attention to what the demon can deduce. The demon has already performed the infinite supertask of fixing all the finitary facts. Surely, it can then step back and observe that it made all those true. This appears to be something the demon can determine a priori from the base physical facts. If it is not to be counted as physical fact, then there is something the demon knows, indeed can come to know a priori, which is not a base fact—which is not on the list of the demon's tasks. Given this close, a priori connection with the base facts, in what way is such a fact not itself base fact? Why is the summation 'all of them' for all the physical facts not itself a physical fact? At some point, to avoid disaster, we will have to say it is not.[20]

We might attempt to explain how such facts could fail to be physical facts by proposing that we are asking too much of the demon. Perhaps it is only able to recognize physical facts as it needs them, and cannot make sense of 'all the facts' in any way but that in which someone dropped in the middle of a forest can talk about 'all the trees'. This seems unsatisfactory. The person in the forest is indeed ignorant of some perfectly ordinary fact—how many trees there are around her. We cannot say likewise of the demon. There cannot be an ordinary physical fact of which it is ignorant. We can only claim that what the demon is ignorant of is not a physical fact, but a kind of super-physical fact. We are thus right back to the hard question, of explaining how 'all the physical facts' can fail to be a physical fact.

The physicalist faces a hard choice. She must cut off infinitary logic at some appropriate $\mathcal{L}_A$, to avoid undermining supervenience. Yet she will thus at some point have to insist that some fact of the 'all of these physical facts' kind is itself a non-physical fact. (Of course, the problem is not restricted to physicalism. The same goes for anyone who wants to hold there is a fundamental class of properties on which other properties supervene.) The puzzle is to explain why any such choice can be right.

I do not want to claim that there is no way to solve this puzzle. How it might be solved will depend on the details of one's philosophical position, so I shall not speculate on which ways will be found better or worse. I only wish to point out that there is a problem here, and that attention to logic shows that is a hard one. Though it has the same source as those raised in the preceding sections, this puzzle is not an analog or refinement of a problem that appeared with respect to finitary closure. Attention to infinitary logic really does pose some new and difficult metaphysical problems for supervenience.

Appendix

Here I sketch a proof of the claim that that many nice fragments of infinitary logic cannot contain a sentence that is equivalent to the conjunction of all true sentences of the fragment. I shall assume familiarity with the basic works of Barwise (1975) and Moschovakis (1974).

Actually, I shall show that many nice fragments cannot contain their own truth predicates. This suffices, as these fragments contain enough syntax that if they contained sentences equivalent to the conjunction of all their true sentences, they would also contain truth predicates.

So, what I shall do is show that in reasonably nice cases, the usual Tarski undefinability of truth result can be extended to infinitary languages. The main difficulty is that to talk of truth predicates at all, we need some syntax coding. For infinitary languages, this is slightly tricky. (Gödel's original result that syntax can be coded in arithmetic relies on some particular properties of ω.) I shall here restrict my attention to those infinitary fragments which behave enough like finitary ones to provide for relatively easy syntax coding.

To this end, let us consider a countable structure $\mathfrak{M}$ with countable language $\mathcal{L}$. Let us assume that $\mathfrak{M}$ is acceptable in the sense of Moschovakis (1974), and furthermore that there is a hyperelementary coding of the the domain M of $\mathfrak{M}$ in the coding system $\mathbb{N}^{\mathfrak{M}}$. There are many such structures, including $\mathbb{N}$ itself, and structures of the form $\langle \mathbb{N}, R \rangle$ which expand $\mathbb{N}$. The languages I shall consider here are the admissible fragments $\mathcal{L}_{HYP_{\mathfrak{M}}}$ for such $\mathfrak{M}$.

We need a notation system that allows us to define the syntax of $\mathcal{L}_{HYP_{\mathfrak{M}}}$ in $\mathcal{L}_{HYP_{\mathfrak{M}}}$. We may begin by observing that for such $\mathfrak{M}$, $HYP_{\mathfrak{M}}$ is projectible into $\mathfrak{M}$ (Barwise 1975, Theorem VI.4.12). Inspecting the proof of this theorem, we can see that the availability of a hyperelementary coding of M in $\mathbb{N}^{\mathfrak{M}}$ allows us to assume that the notations are all in Seq, the elementary collection of $\mathbb{N}^{\mathfrak{M}}$-sequences.

It is well-known that the syntax of $\mathcal{L}_{HYP_{\mathfrak{M}}}$ is Δ on $HYP_{\mathfrak{M}}$ in the language $\mathcal{L}^* = (\mathcal{L}, \in)$. Together with the projectibility of $HYP_{\mathfrak{M}}$, this shows that each syntactic relation is a Δ relation on M. Furthermore, as each notation is in Seq, each syntactic relation is a Δ relation on the elementary Seq. Hence, by Δ-separation, each syntactic relation is a relation on M in $HYP_{\mathfrak{M}}$. Now, we can appeal to a general theorem that any relation on M in $HYP_{\mathfrak{M}}$ is $\mathcal{L}_{HYP_{\mathfrak{M}}}$-definable (Barwise 1975, Corollary IV.3.5). This gives us the $\mathcal{L}_{HYP_{\mathfrak{M}}}$-definability of the syntax of $\mathcal{L}_{HYP_{\mathfrak{M}}}$.

Once we have syntax coding in hand, we can just repeat the usual proof of the undefinability of truth. Suppose, for contradiction, that $\mathcal{L}_{HYP_{\mathfrak{M}}}$ contains a predicate Tr which is true just of true sentences of $\mathcal{L}_{HYP_{\mathfrak{M}}}$, i.e.:

$$\mathfrak{M} \models Tr(\ulcorner \phi \urcorner) \leftrightarrow \phi.$$

Given that the syntax of $\mathcal{L}_{HYP_{\mathfrak{M}}}$ is $\mathcal{L}_{HYP_{\mathfrak{M}}}$-definable, we can apply the usual proof of the diagonal lemma (as in, e.g., Boolos and Jeffrey 1989) to $\mathcal{L}_{HYP_{\mathfrak{M}}}$ to find a sentence λ such that:

$$\mathfrak{M} \models \neg Tr(\ulcorner \lambda \urcorner) \leftrightarrow \lambda.$$

Putting the two together, we have:

$$\mathfrak{M} \models Tr(\ulcorner \lambda \urcorner) \leftrightarrow \neg Tr(\ulcorner \lambda \urcorner),$$

the familiar contradiction.[21]

Notes

[1]For what happens in this paper, it will not be important just what the status of properties is. Nearly everything said here would hold if we started with predicates rather than properties. The early Hellman and Thompson (1975) describes supervenience in terms of facts rather than properties, but in developing their view, they really work with predicates applied to objects, so the difference is not important to this discussion.

[2]The variety of supervenience relations was discussed by Teller (1983b), and then Kim (1984), which gave us the familiar list of weak, strong, and global supervenience. The distinct contributions of covariation and modalization are nicely brought out by McLaughlin (1995) and Bacon (1995), which calls a version of my covariation 'protosupervenience'. Some scrutiny of the modal notions involved in supervenience may be found in Stalnaker (1996) and Wedgwood (2000). The question of whether supervenience should be 'individualistic'—considering only properties of individual objects—is raised by Petrie (1987) and discussed in depth by Post (1995). The idea of regional supervenience is due to Horgan (1993).

[3]Among places where one sees the use of Boolean closure assumptions are the many works of Kim (e.g. Kim 1984, 1990), Bacon (1986), Paull and Sider (1992), and Stalnaker (1996).

[4]For specific applications of supervenience, the presence of relations in the base class may well raise some subtle issues. For instance, what exactly counts as the physical relations an object stands in, and whether they do or do not suffice for a desired supervenience relation, can raise all sorts of questions. Contingent existence can further complicate matters. For some discussion, see Kim (1987, 1993a), Petrie (1987), and Post (1995). The 'Humean supervenience' of Lewis (e.g. Lewis 1986), explicitly supposes the base to include spatio-temporal relations, but no other relational properties.

These subtleties do not matter for my point here. All I am saying is if for some reason we settle on a supervenience base which contains some relations, and we also wish to assume

Boolean closure, we should in fact be assuming full $\mathcal{L}$-closure.

[5]The 'sentences' of an infinitary language like $\mathcal{L}_{\infty\omega}$ are mathematical abstractions, usually defined in set theory. If you like, you can think of the disjunction $\bigvee \Phi$ as being nothing other than the set $\langle \vee, \Phi \rangle$. It does not matter that in crucial ways these sets fail to resemble the sentences of natural languages. The sentences of more familiar finite formal languages, like first order logic, also fail to resemble the sentences of natural languages in crucial ways. Aside from not following the syntactic rules of natural languages, as Sylvain Bromberger delights in pointing out, these formalisms do not have a phonology, and so cannot be uttered in the normal sense.

[6]In a much-cited passage, Teller (1983a, p. 58) notes that the infinite disjunctions involved in Kim's argument will be "awfully fat."

[7]Whether these steps works for other supervenience relations is a question of some delicacy. Jackson (1998), for instance, gives a version of it for the supervenience of ethical on physical properties. Jackson works primarily with a global supervenience relation, rather than the strong supervenience of Kim's argument, but he does note that his version of the argument relies on special features of the relation between ethical and physical properties. There is some debate over whether or not global supervenience relations support the reduction argument generally. See Petrie (1987) and Grimes (1995). In the background here is the long-running dispute over the relation of global to strong supervenience. See, in addition to Petrie's paper, Kim (1984), Paull and Sider (1992), and Stalnaker (1996).

[8]This result is known as the Scott isomorphism theorem, due to Dana Scott, generalized by Chang. A proof may be found in Barwise (1975). There are some technical subtleties about the notion of 'fully describing' a structure. Scott sentences characterize countable structures up to isomorphism. If we look at uncountable structures, we get only what is called 'partial isomorphism'. Partial isomorphism is slightly weaker than isomorphism, but a theorem of Karp shows partial isomorphism to be equivalent to the relation of making all the same

sentences of $\mathcal{L}_{\infty\omega}$ true. This is adequate to ensure 'sameness of word', for our purposes at least.

[9]The original argument is due to Bacon (1986). Critical responses were offered by Van Cleve (1990) and Oddie and Tichý (1990). Bacon (1990, 1995) replies.

[10]I find this issue somewhat puzzling. Where, in naming a physical object, did we somehow step outside the physical? Where did we invoke such a strong metaphysical tool as an essence? Not, surely, in the ostension and baptism of the object. On the other hand, it may be replied, in introducing a name, a rigid term, we entangle ourselves in the difficult issues of transworld identity. Perhaps in doing so, we tacitly rely on some strong metaphysics?

[11]There is one technical matter that bears mention here. Worlds are taken to be $\mathcal{L}$-structures, where the $\mathcal{L}$-vocabulary is the vocabulary of $\mathcal{B}$. Technically, this is the natural way to proceed. However, it may impose a kind of global supervenience constraint, as the only differences in worlds to which we can appeal are $\mathcal{B}$-differences. The matter is actually a bit tricky, as whether this amounts to global supervenience depends on just what global supervenience has in mind by $\mathcal{B}$-differences. If it is just a matter of what finitary $\mathcal{L}$-sentences worlds make true, then we have lots of distinct worlds that have no $\mathcal{B}$-differences. Any elementarily equivalent non-isomorphic structures will be examples. If on the other hand, if it is a matter of what $\mathcal{L}_{\infty\omega}$-sentences worlds make true, then we cannot violate it with the techniques I have discussed here.

This may indeed restrict the range of properties we can add to those which satisfy a global supervenience constraint, and it may likewise limit the amount of resplicing that can be done. However, I do not think this matters very much. Even if there is a restriction, there are still lots of wild resplicing that can be done, and it does not affect the issue of adding rigid properties—sets—at all. As I mentioned, all I really need here is a preponderance of evidence that $\mathcal{L}_{\infty\omega}$-closure has objectionable consequences, and we have plenty of that whether the restriction is significant or not.

Incidentally, the arbitrary property construction does does provide an argument that under the hypothesis of $\mathcal{L}_{\infty\omega}$-closure, strong and global supervenience coincide. Stalnaker (1996) gave an argument for this relying on infinite quantifier prefixes, at a minimum $\mathcal{L}_{\omega_1\omega_1}$. My approach avoids this in favor of the much more tractable $\mathcal{L}_{\infty\omega}$.

[12]This is a fact well-known to specialists. The best examples come from way infinitary logic can describe ordinals.

[13]Experience has also shown that restrictions based on the size of Φ tend not to be helpful for the problems we face here. Most of the difficulties of Section (2) can appear with even the smallest infinitary language in terms of cardinality, $\mathcal{L}_{\omega_1\omega}$, so long as we restrict attention to a countable set of countable structures.

[14]For those who care about the details: The most important conditions of admissibility are Δ_0-separation and Σ-replacement. The other requirements on admissibles are that they satisfy some of the ordinary axioms of set theory (in unrestricted form): extensionality, foundation, pairing, and union. See Barwise (1975) for an extended discussion.

For specialists: A^+ is basically the sets that are hyperelementary over A.

[15]Again, for those who care: A fragment is a collection of formulas closed under finitary logical operations, as well as subformulas and substitution of terms. An admissible fragment is the intersection of an admissible with $\mathcal{L}_{\infty\omega}$. See Barwise (1975) for details.

[16]The case of Scott sentences is technically quite subtle, and really goes beyond the scope of this paper. However, to gesture at some important results: Nadel (1974) shows that for admissible A and structure $\mathfrak{M} \in A$, $\sigma_{\mathfrak{M}} \in A^+$. It need not be in A. Nadel also shows that for $\mathfrak{M}, \mathfrak{N} \in A$, $\mathfrak{M} \equiv \mathfrak{N}(\mathcal{L}_{\infty\omega})$ iff $\mathfrak{M} \equiv \mathfrak{N}(\mathcal{L}_A)$. Hence, $\mathcal{L}_A$ suffices to characterize structures in A up to partial isomorphism.

[17]The set of hereditarily finite sets is admissible, so we can always find an admissible ruling out any infinite set. On the other hand, the A^+ construction can be used to build an admissible containing a given set.

[18]Either the conjunction or disjunction of all the true sentences of a given language works like a truth predicate for that language. The idea that the relation of supervenient to base properties bears some resemblance to that of truth predicate to the true sentences of a language was, to my knowledge, first considered by Davidson (1970).

[19]Jackson (1998, p. 26) entertains the idea that the collection of physical facts must contain a "stop" or "that is all" clause, and insists it is of a "purely physical character." He has in mind the problem for physicalists of ruling out worlds that are physically like ours but also contain non-physical stuff. Though our settings are somewhat different, I too am arguing for something like a stop clause. However, in my setting, granting its purely physical character undermines its efficacy.

[20]In making such deductions, the demon might perform something much like what Horgan (1983) calls 'cosmic hermeneutics'. Horgan's particular worry is whether the demon can deduce facts expressible in the vocabulary of the supervenient class, which raises questions about the a posteriori status of facts about meaning. Some, notably Jackson (1998) and Chalmers (1996), endorse the prospect of cosmic hermeneutics. For critical discussion, see Byrne (1999).

There are other problems for the demon if infinitary logics are around. For one thing, once we leave countable admissibles, we get failures of compactness, often very bad ones. But then, even if the demon is correct in some conclusion, it may be unclear in what sense it has proved the conclusion.

[21]Andrew Botterell and Daniel Stoljar convinced me, or more forced me, to write this paper. I am also extremely grateful to Alex Byrne, Ned Hall, Jim Pryor, Susanna Siegel, Zoltán Gendler Szabó, Judith Thomson, Ralph Wedgwood, Steve Yablo, and two anonymous referees for many helpful comments and discussion.

References

Armstrong, D. M. (1978). A Theory of Universals: Universals and Scientific Realism Volume II, Cambridge University Press, Cambridge.

Bacon, John (1986). "Supervenience, Necessary Coextension, and Reducibility," Philosophical Studies **49**: 163–176.

Bacon, John (1990). "Van Cleve Versus Closure," Philosophical Studies **58**: 239–242.

Bacon, John (1995). "Weak Supervenience Supervenes," in Elias E. Savellos and Ümit D. Yalçin (eds.), Supervenience: New Essays, Cambridge University Press, Cambridge, pp. 101–109.

Barwise, Jon (1975). Admissible Sets and Structures, Springer-Verlag, Berlin.

Boolos, George and Jeffrey, Richard (1989). Computability and Logic, 3rd edn., Cambridge University Press, Cambridge.

Byrne, Alex (1999). "Cosmic Hermeneutics," in James E. Tomberlin (ed.), Epistemology, vol. 13 of Philosophical Perspectives, Blackwell Publishers, Oxford, pp. 347–383.

Chalmers, David J. (1996). The Conscious Mind, Oxford University Press, Oxford.

Davidson, Donald (1970). "Mental Events," in L. Forest and J. Swanson (eds.), Experience and Theory, University of Massachusetts Press, Amherst, pp. 79–101. Reprinted in Davidson (1980).

Davidson, Donald (1980). Essays on Actions and Events, Oxford University Press, Oxford.

Grimes, Thomas R. (1995). "The Tweedledum and Tweedledee of Supervenience," in Elias E. Savellos and Ümit D. Yalçin (eds.), Supervenience: New Essays, Cambridge University Press, Cambridge, pp. 110–123.

Hellman, Geoffrey Paul and Thompson, Frank Wilson (1975). "Physicalism: Ontology, Determination, and Reduction," Journal of Philosophy **72**: 551–564.

Horgan, Terence (1983). "Supervenience and Cosmic Hermeneutics," in Terence Horgan (ed.), Spindel Conference 1983: Supervenience, vol. 22 of Southern Journal of Philosophy Supplements, pp. 19–38.

Horgan, Terence (1993). "From Supervenience to Superdupervenience: Meeting the Demands of the Material World," Mind **102**: 555–586.

Jackson, Frank (1998). From Metaphysics to Ethics, Oxford University Press, Oxford.

Kim, Jaegwon (1984). "Concepts of Supervenience," Philosophy and Phenomenological Research **45**: 153–176. Reprinted in Kim (1993b).

Kim, Jaegwon (1987). "'Strong' and 'Global' Supervenience Revisited," Philosophy and Phenomenological Research **48**: 315–326. Reprinted in Kim (1993b).

Kim, Jaegwon (1990). "Supervenience As a Philosophical Concept," Metaphilosophy **21**: 1–27. Reprinted in Kim (1993b).

Kim, Jaegwon (1993a). "Postscript on Supervenience," in Supervenience and Mind, Cambridge University Press, Cambridge, pp. 161–171.

Kim, Jaegwon (1993b). Supervenience and Mind, Cambridge University Press, Cambridge.

Lewis, David (1986). Philosophical Papers, vol. II, Oxford University Press, Oxford.

McLaughlin, Brian P. (1995). "Varieties of Supervenience," in Elias E. Savellos and Ümit D. Yalçin (eds.), Supervenience: New Essays, Cambridge University Press, Cambridge, pp. 16–59.

Moschovakis, Yiannis N. (1974). Elementary Induction on Abstract Structures, North-Holland, Amsterdam.

Nadel, Mark (1974). "Scott Sentences and Admissible Sets," Annals of Mathematical Logic **7**: 267–294.

Oddie, Graham and Tichý, Pavel (1990). "Resplicing Properties in the Supervenience Base," Philosophical Studies **58**: 259–269.

Paull, R. Cranston and Sider, Theodore R. (1992). "In Defense of Global Supervenience," Philosophy and Phenomenological Research **52**: 833–854.

Petrie, Bradford (1987). "Global Supervenience and Reduction," Philosophy and Phenomenological Research **48**: 119–130.

Post, John F. (1983). "Comment on Teller," in Terence Horgan (ed.), Spindel Conference 1983: Supervenience, vol. 22 of Southern Journal of Philosophy Supplements, pp. 163–167.

Post, John F. (1995). ""Global" Supervenient Determination: Too Permissive?" in Elias E. Savellos and Ümit D. Yalçin (eds.), Supervenience: New Essays, Cambridge University Press, Cambridge, pp. 73–100.

Stalnaker, Robert (1996). "Varieties of Supervenience," in James E. Tomberlin (ed.), Metaphysics, vol. 10 of Philosophical Perspectives, Blackwell Publishers, Oxford, pp. 221–241.

Teller, Paul (1983a). "Comments on Kim's Paper," in Terence Horgan (ed.), Spindel Conference 1983: Supervenience, vol. 22 of Southern Journal of Philosophy Supplements, pp. 57–61.

Teller, Paul (1983b). "A Poor Man's Guide to Supervenience and Determination," in Terence Horgan (ed.), Spindel Conference 1983: Supervenience, vol. 22 of Southern Journal of Philosophy Supplements, pp. 137–162.

Van Cleve, James (1990). "Supervenience and Closure," Philosophical Studies **58**: 225–238.

Wedgwood, Ralph (2000). "The Price of Non-Reductive Physicalism," Noûs **34**: 400–421.